Track Extraction with Hidden Reciprocal Chain Models

George Stamatescu, *Student Member, IEEE,* Langford B White, *Senior Member, IEEE,* Riley Bruce-Doust*

arXiv:1605.04046v1 [cs.CV] 13 May 2016

***Abstract*—This paper develops Bayesian track extraction algorithms for targets modelled as hidden reciprocal chains (HRC). HRC are a class of finite-state random process models that generalise the familiar hidden Markov chains (HMC). HRC are able to model the "intention" of a target to proceed from a given origin to a destination, behaviour which cannot be properly captured by a HMC. While Bayesian estimation problems for HRC have previously been studied, this paper focusses principally on the problem of track extraction, of which the primary task is confirming target existence in a set of detections obtained from thresholding sensor measurements. Simulation examples are presented which show that the additional model information contained in a HRC improves detection performance when compared to HMC models.**

I. Introduction

1) Problem Description: Track extraction is a statistical hypothesis testing technique that tests for the presence of a dynamic target in a set of detections of uncertain origin that are obtained from a detector applying a low threshold to sensor measurements [1], [2]. While standard multiple hypothesis tracking (MHT) algorithms confirm or delete tracks based on the likelihood of individual candidate tracks [3], track extraction algorithms use the entire set of detections to first decide whether a target exists, with the likelihood supplied by applying a tracker with an underlying model of target dynamics. It is therefore also a detection process, with quantities such as detection probability and false alarm rate, but works on a higher level of abstraction than a sensor detector, leading to tracks rather than potentially unrelated point estimates. Classically, tracking and track extraction algorithms apply the Markov assumption to the target dynamics. At finer time scales the assumption is generally appropriate, and well-known Markov based tracking algorithms are discussed in the literature [3],[4]. This paper is motivated by tracking applications with coarser time scales such as [5]-[8] where the Markov assumption is less valid generally because targets have a defined origin and destination. We will call such targets *source-destination aware*, and as in [5]-[8] we work with finite-state processes because the arbitrary source-destination pair relationships are better captured by a cellular state space. In contrast to [6]-[8], which largely focus on applications and only model a target's destination implicitly or via an ad hoc approach, we examine explicitly whether the use of dynamic target models which incorporate such information can improve track extraction.

George Stamatescu, Langford B White and Riley Bruce-Doust are with the School of Electrical and Electronic Engineering, The University of Adelaide, North Tce, Adelaide 5005, Australia. {George.Stamatescu, Lang.White, Riley.Bruce-Doust}@adelaide.edu.au

2) Related Work: From a modelling point of view, being source-destination aware means that the initial and final target states have a specified joint probability distribution and that target dynamics are anticipative, or non-causal, reflecting the intention of the target to move towards its destination. There have been a range of approaches to incorporating available *a priori* information about the future within the class of Markov models, both at the estimation stage and in the target dynamics model itself. Both methods amount to "back propagating" the influence of a future state value or distribution by appropriate conditioning. For example in [9],[10], *a priori* destination information is incorporated via a corrective term into the state update of continuous-time Gauss-Markov processes, and in [11]-[13] specifying an *a priori distribution* for a future time of the state process creates a new Markovian process referred to as a Schrödinger bridge. What is however not able to be modelled within the Markovian class is the future information that is conditional on the past state of the system (the intention to proceed between a source and associated destination). The injection of a relationship between future and past states other than that induced by the evolution of the globally naive first order Markovian dynamics, adds a probabilistic dependence that raises the class of model to what is known as the class of *reciprocal processes*. Reciprocal processes were studied in detail by Jamison [14] in a general setting, and subsequently by Levy *et al* [15] who considered the realisation and state estimation problem for the Gaussian discrete parameter case. There are many other related works which are summarised in [16], to which to the interested reader is referred for more background.

Reciprocal chains (RCs) are finite-state reciprocal processes, and hidden reciprocal chains (HRCs) are stochastic processes generated from an RC via some noisy and/or incomplete observation mechanism, analogous to hidden Markov chains (HMC). As stated we work with finite-state processes primarily because we are concerned with targets with arbitrary source-destination pair relationships and the precision of point estimates produced by continuous Gauss-Markov processes are not relevant. Finite states also allow for state space constraints, such as "forbidden areas" for the target or restrictions to roads as in [5]-[8], and more abstract target attribute states. We thus recount the present state of development in inference of HRCs; un-normalised optimal filters/smoothers for HRC, derived using a Bayesian approach, were presented in [16]. Normalised filters and smoothers were developed in [17],

which also considered a generalisation of HRC to incorporate "waypoints" on the target trajectory. Maximum likelihood estimation of state sequences (MLSE) for HRC was presented in [18]. In [19],[20] an approach called destination aware tracking based on HRC was first proposed, but the RC models tested, being pinned to a single destination, in fact remained in the Markov class. Thus the tracking performance of non-causal HRCs has not been studied in detail. The present paper is an extension of the work presented in [21], which provided an intuitive description of how HRCs may be applied in a target tracking problem. The filtering algorithms derived in this paper proceed in the manner of the Bayesian decomposition of a reciprocal chain into a collection of end-pinned bridge processes that was established in [16].

3) Contributions: The key contributions of this paper are in answering the questions of whether modelling the source-destination awareness of a target can improve track extraction, and to what extent it is improved as a function of the parameters of a HRC, which model source-destination awareness. The track extraction setting to which we introduce the new model, and compare to existing models, is similar to the setup in [1], with three key differences. Firstly, we perform optimal Bayesian estimation over finite states, rather than a continuous space. Secondly we operate on a fixed interval rather than in an online context, and thirdly the likelihood ratio test here has a single threshold, though it could be formulated as the sequential probability ratio test of [1] which has two thresholds. To recreate the setting, we include two types of uncertainty, false alarms - detections that do not originate from the target (ie. from 'clutter), and an arbitrary 'sensor' noise which degrades position estimates. This could noise could include the effect of quantisation of sensor measurements associated with the finite state space. In section V of [1], signal-strength information is used and improves the track extraction performance, but here we remove any explicit notion of SNR and signal amplitude to emphasise the role of the model in differentiating clutter from target detections. Instead we treat the association between observation and the source of the detection probabilistically, via an observation likelihood function and *a priori* probabilities over false alarms. Observation models for two representative track extraction regimes are developed and tested. In each there is at most one target during the tracking interval. In the first we process one sensor detection at each time increment, with a given fixed *a priori* probability it is a false alarm. In the second scenario we generalise this to allow multiple detections at each time, with at most one due to the target, a scenario also typical in visual tracking and which is very similar to the observation model of [1]. A dynamic model of the target is constructed which incorporates target dynamics, including the global source-destination attributes. The likelihood associated with a given set of detections is evaluated using the new class of normalised hidden reciprocal chain (HRC) filters. This likelihood forms the basis for a target extraction algorithm. With most finite data non-linear inference problems, it is infeasible to obtain *a priori* performance metrics, so we are led to the use of numerical simulations to study the performance of the new algorithms proposed here.

There has been a steady shift across tracking application domains towards modelling more complex target behaviours, for example in user support applications such as location based services (eg. [5],[6]), and in tracking scenarios where analysts' decision making can benefit from models that suit questions of higher level inference, such as predicting destination (eg. [19]). The contributions of this paper are relevant to these application areas as they describe, via numerical studies, the tracking benefit that stands to be gained by incorporating a simple notion of intent. Alternative approaches [5]-[8] have largely assumed rather than demonstrated that modelling intent improves tracking, which this paper focuses on explicitly within the track extraction setting.

4) Paper Outline: The layout of the paper is as follows: in Section II, we define HRC and summarise the pinned Markov process construction of an HRC such as given in [6]. The observation models are defined in Section III before the normalised optimal filters are derived in Section IV, and the corresponding detectors in section V. Numerical simulations are presented in Section VI, before we conclude in Section VII with a discussion of future research directions.

II. Hidden Reciprocal Chain Models

In this section, we define a reciprocal chain (RC) and summarise the Markov bridge (MB) construction of a RC as outlined in [16]. We then describe how source-destination awareness is encoded in the RC model, and how this awareness can be approximated with a Schrödinger bridge. We then define a HRC as a natural generalisation of a RC in the same way that a HMC generalises a Markov chain (MC). The HRC construction is used to derive the optimal HRC filters in section IV.

A. The Markov Bridge Construction of a RC

Consider a random process $\{X_t\}$ indexed by $t \in \{0, 1, \ldots, T\}$ for some fixed integer $T \geq 2$. The index t will denote an index on a set of event epochs. At each epoch t, the random variable X_t takes a value on a finite state space $\mathcal{S} = \{1, \ldots, N\}$, where $N \geq 2$ is a fixed integer. The process $\{X_t\}$ is said to be *reciprocal* [14], if

$$\mathbf{P}\left\{X_t | X_s, \forall\, s \neq t\right\} \quad = \quad \mathbf{P}\left\{X_t | X_{t-1}, X_{t+1}\right\} \ , \qquad (1)$$

for each $t = 1, \ldots, T-1$. Thus X_t is conditionally independent of $X_0, \ldots, X_{t-2}, X_{t+2}, \ldots, X_T$ given its neighbours X_{t-1} and X_{t+1}. The reciprocal model is specified by the set of three-point transition functions (1) together with a given joint distribution on the end points $\mathbf{P}\{X_1, X_T\}$. A reciprocal process is not, in general, a Markov process, however all Markov processes are reciprocal [14]. Denote the three-point transition functions in (1) by

$$Q_{i,j,\ell}(t) = \mathbf{P}\left\{X_t = j | X_{t-1} = i, X_{t+1} = \ell\right\}), \qquad (2)$$

for $i, j, \ell \in \mathcal{S}$, $t = 1, \ldots, T-1$, and the end-points distribution given by

$$\mathbf{P}\left\{X_0 = i, X_T = j\right\} \quad = \quad \Pi_{i,j}, \quad i, j \in \mathcal{S} \ . \qquad (3)$$

Pinning the end point of a RC generates a Markov bridge, which we consider to be a Markov process with the end point X_T fixed to a specified value, together with the initial state distribution, obtained from (3), conditioning on the value X_T takes. The term bridge generally refers to a process with specified values at both start and end times, however provided the end is reachable one may define it with an initial distrubtion rather than fixed initial state.

So for the finite (N) state case, a RC can be regarded as N Markov bridges, one corresponding to each of the possible final states taken by X_T. When we consider the joint distribution of the states of a RC and using direct Bayes' conditioning and (1), the relevance of Markov bridges becomes apparent,

$$\mathbf{P}\left\{X_0,\ldots,X_T\right\} = \mathbf{P}\left\{X_1,X_T\right\}\left(\prod_{t=2}^{T-1}\mathbf{P}\left\{X_t|X_{t-1},X_T\right\}\right). \tag{4}$$

The terms contained within the product in (4) are precisely the state transitions for a MB pinned at X_T, and depend only on the three-point transitions (2). Based on the properties of a RC a backwards recursion was obtained in [16] which fully specifies the set of N MB transitions for $t = T-2, T-3, \ldots, 1$ via

$$\begin{aligned} B^k_{i,j}(t) &= \mathbf{P}\left\{X_{t+1}=j|X_t=i,X_T=k\right\} \\ &= \frac{Q_{i,j,\ell}(t+1)}{B^k_{j,\ell}(t+1)}\left(\sum_{m=1}^{N}\frac{Q_{i,m,\ell}(t+1)}{B^k_{m,\ell}(t+1)}\right)^{-1}, \end{aligned} \tag{5}$$

the last term on the right being the normalisation constant. Initialisation is with $B^k_{i,j}(T-2) = Q_{i,j,k}(T-1)$. The quantity on the right hand side of equation (5) is independent of the index ℓ. Now since the terms $B^k_{j,.}(t+1)$ form a probability distribution (i.e. they are non-negative and sum to unity) for each i,k and t, there is at least one index ℓ for which $B^k_{j,\ell}(t+1)$ is non-zero. So this index may be selected on the rhs of (5), and thus (5) is well defined. From a numerical perspective, it may thus be appropriate to choose a value of ℓ (for each j,k,t) which maximises the value of $B^k_{j,\ell}(t+1)$ when performing the calculations in (5). Determining the set of MBs from the specified three point transitions requires complexity of $O(N^3T)$, i.e. $O(N^2T)$ for each MB. The Markov bridge with final state k has an initial probability distribution π^k_i given by the conditional distribution

$$\pi^k_i = \mathbf{P}\left\{X_0=i|X_T=k\right\} \;=\; \frac{\Pi_{i,k}}{\sum_{j=1}^{N}\Pi_{j,k}}\,. \tag{6}$$

Thus we have that any RC may be uniquely specified by the finite set of Markov bridges with probability transition matrices given by (5) and initial distributions (6). As we will see below, a more intuitive description of RCs can be found via the construction of a reciprocal process from a Markov process, which we will refer to as the *base process*.

B. Encoding Source Destination Awareness

As stated in the introduction, this paper compares the performance of a tracker derived from a reciprocal model to a tracker whose target model includes future information but remains within the Markovian class. Specifically, we will compare to the Schrödinger bridge (SB) [11]–[13]. In this subsection we describe both the RC and SB approaches to encoding future information into a Markov process representing a target model with given Markov transition probability matrices, referred to as the base process. We then present an example of a non-Markov RC specified by Π.

Reciprocal chains from a base process:

This construction was established in [14] and will be used here, as it was in [16], for numerical simulations in Section VI. Let A denote the time-homogeneous transition probability matrix for a N state Markov chain Z_t, which we call the base process, where the entries of the matrix are $A_{i,j} = \mathbf{P}\left\{Z_t = j|Z_t = i\right\}$. A reciprocal process can be constructed from the base process with three point transitions given by [14]

$$Q_{i,j,\ell}(t) = \frac{A_{i,j}A_{j,k}}{\sum_{\ell=1}^{N}A_{i,\ell}A_{\ell,k}}\,,$$

for $t = 1,\ldots,T-2$. As described in [16], The corresponding MB transitions (5), are now given by

$$B^k_{i,j}(t) = \frac{A_{i,j}(A^{T-(t+1)})_{j,k}}{(A^{T-t})_{i,k}}\,, \tag{7}$$

for $t = 0,\ldots,T-2$. As described in [16], the process for generating a sample path of this RC is to draw the initial and final points X_0 and X_T from Π, which specifies the MB transitions corresponding to $X_T = k$. The sample path is then constructed in the standard way for a MC, starting from X_0 using the transitions (7).

The end-points distribution obviously encodes the target's source-destination awareness, as well as determining whether a particular RC is Markov or not. An example of a non-Markov destination awareness which we call *loitering behaviour* where a target returns to its origin by some final time T, can be encoded via

$$\mathbf{P}_{Loitering}\left\{X_0=i,X_T=j\right\} = \begin{cases} p_r(i), & \text{if } i=j, \\ 0, & \text{otherwise.}\end{cases} \tag{8}$$

where $p_r(i)$ is the probability of starting and returning to the origin i, and $\sum_{i=1}^{N}p_r(i)=1$. If there is a uniform probability on loitering anywhere in a statespace with N states, the end-points distribution for this loitering reciprocal chain (LRC) model can be written in matrix form;

$$\Pi_{LRC} \;=\; \begin{pmatrix} \frac{1}{N} & 0 & \cdots & 0 \\ 0 & \frac{1}{N} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{1}{N} \end{pmatrix}$$

We can see clearly how "source-destination pairs" of a target are encoded via Π.

Schrödinger bridges from a base process:

The SB is the time-inhomogeneous process that attains a specified marginal distribution at its final time and which has transitions, i.e. 'dynamics', closest to (in Kullback-Liebler

information divergence sense) the specified *a priori* homogeneous dynamics of the base process (see [13] for a finite-state proof). We thus recount the Schrödinger bridge transition matrix construction. The following originates from E. Schrödinger's idea, was formalised in probabilistic terms by Jamison [11], and was recently translated into the discrete-state setting by Pavon *et al* [12], [13]. Let A denote the homogeneous state transition matrix of a base process Z_t, as before, and let X_t be the SB of Z_t with marginal distributions π_0 and π_T on X_0 and X_T respectively. Let λ_0, λ_T be the solutions of

$$\pi_T = \lambda_T \circ \lambda_0 A^T$$
$$\pi_0 = \lambda_0 \circ \lambda_T A'^T$$

where $\circ$ is the element-wise product. (Existence and uniqueness is proved in [13].) Define the positive row vectors $\psi_t = \lambda_T A'^{(T-t)}$. Then the transitions of the SB are given by

$$S_{i,j}(t) = \mathbf{P}\left\{X_{t+1} = j | X_t = i\right\} = A_{i,j}\frac{\psi_{t+1}(j)}{\psi_t(i)} \quad (9)$$

We can attempt to approximate a RC with a SB by setting the marginals π_0 and π_T to be precisely those of the RC's endpoints distribution Π. As we will see in section VI, although SBs are a way to encode future information, they are unable to capture the specific source-destination awareness that non-Markov reciprocal processes can, as there cannot be an arbitrary statistical relationship between the start and endpoints of a target.

C. Hidden Reciprocal Chains (HRCs)

Suppose that the RC $\mathcal{X} = \{X_0, , \ldots, X_T\}$ is observed via the observation process $\mathcal{Y} = \{Y_0, \ldots, Y_T\}$. Assume that the observation at epoch t given the state X_t is conditionally dependent of X_τ and Y_τ, $\tau \neq t$. This conditional independence implies that

$$\mathbf{P}\left\{Y_0, \ldots, Y_T | X_0, \ldots, X_T\right\} \quad = \quad \prod_{t=0}^{T} \mathbf{P}\left\{Y_t | X_t\right\}. \quad (10)$$

The process $\mathcal{Y}$ is called a hidden reciprocal chain (HRC) because the property (10) is analogous to the usual assumption made for hidden Markov chains. Let

$$C_i(t) \quad = \quad \mathbf{P}\left\{Y_t | X_t = i\right\} \ , \quad (11)$$

denote the (conditional) observation likelihoods of the HRC. The observations may be either discrete or continuous random variables defined on an appropriate probability space. In the continuous observation case, each $C_i(t)$ would denote a probability density function, parametrised by a finite set of parameters. In the discrete (finite) observation case, the $C_i(t)$ would represent a probability mass function with a finite number of mass points. The mixed continuous/discrete observation case is of course, also possible. This observation process models targets in an uncluttered environment, in the following section we present two models of the observation process when there is clutter present.

III. Observation Models

In this section, we define the two observation models which correspond to the tracking scenarios of interest, that is, a set of detections over an interval obtained by applying a threshold to sensor returns. The first model allows a single detection at each time, that may be a target detection or a false alarm, modelled with position uncertainty. The second model generalises this to allow multiple detections at each time, with at most one detection corresponding to a target. Note that the single observation model could alternatively be interpreted as potentially erroneous output from a fast-time tracker, such as a Kalman filter (see [19] for further details), while the second model is more typical of general tracking scenarios, and similar to that of [1].

A. Single Observation Model

A general threshold based detection process can be constructed by proposing a model for a detector which has two operating states $\mu_t = \{0, 1\}$ at each time t, one for when it detects the target ($\mu_t = 1$), and another for false alarms or non-target detections ($\mu_t = 0$). We denote the *a priori* probabilities over μ_t as

$$\mathbf{P}\left\{\mu_t = 0\right\} = \epsilon$$
$$\mathbf{P}\left\{\mu_t = 1\right\} = 1 - \epsilon$$

where $\epsilon \in [0, 1]$. The observation model (conditional distribution) for this scenario can then be derived as a mixture of the form

$$\begin{aligned}\mathbf{P}\left\{Y_t | X_t = i\right\} &= \sum_{\mu_t=0,1} \mathbf{P}\left\{Y_t, \mu_t | X_t = i\right\} \\ &= \sum_{\mu_t=0,1} \mathbf{P}\left\{\mu_t\right\} \mathbf{P}\left\{Y_t | X_t = i, \mu_t\right\} \\ &= (1-\epsilon)\mathbf{P}\left\{Y_t | X_t = i, \mu_t = 1\right\} \\ &\qquad + \epsilon \mathbf{P}\left\{Y_t | \mu_t = 0\right\} \quad (12)\end{aligned}$$

We define the first term of Equation (12) as another observation likelihood, given we know the operating mode, as $c_i(t) = \mathbf{P}\left\{Y_t | X_t = i, \mu_t\right\}$. As before $c_i(t)$ can take any form as in definition (11). However, the lower case notation is deliberate as it signifies it is the observation likelihood given the state *and* the operating mode of the sensor.

The final term of (12), represents the likelihood of the observation Y_t given that it is of non-target origin, hence the lack of dependence on the state X_t. While the source of such a detection is unknown, it is reasonable to assume it is generated it in a similar way as a target detection and hence can take the same observation likelihood as $c_i(t)$, given we know the operating mode. Specifically, we define

$$\begin{aligned}\mathbf{P}\left\{Y_t | \mu_t = 0\right\} &= \sum_{j=1}^{N} \mathbf{P}\left\{Y_t | U_t = j, \mu_t = 0\right\} \mathbf{P}\left\{U_t = j\right\} \\ &= \sum_{j=1}^{N} c_j(t) \mathbf{P}\left\{U_t = j\right\}\end{aligned}$$

where U_t is a temporally uncorrelated process over the state-space representing clutter. For brevity, we will refer to the sensor detector dalse alarm rate ϵ, as the clutter rate.

B. Multiple Observation Model

We can generalise the single observation case to M multiple detections at a time t, with at most one target in the system and potentially no target detections at any particular time. Thus we have that the observations at each time t are a list $Y_t = \{y_t^1, \dots y_t^M\}$.

We can define an *association* random vector $\mathbf{a_t}$ which acts as a pointer to the list of observations, associating an observation to a target detection. It is convenient to let $\mathbf{a_t}$ take as values the zero vector or one of the orthonormal basis vectors in $\mathbb{R}^M$,

$$\mathbf{a_t} \in \{\mathbf{0}, \mathbf{e_1}, \mathbf{e_2}, \dots, \mathbf{e_M}\}$$

Hence $\mathbf{a_t}$ has a 1 in at most one position, and zeros elsewhere, with entries corresponding to the entries of the list Y_t. We could view each entry of $\mathbf{a_t}$ as a mode just as in the previous observation model. We denote the *a priori* probability over $\mathbf{a_t}$ as,

$$\begin{aligned} \mathbf{P}\{\mathbf{a_t} = \mathbf{0}\} &= \lambda_0 \\ \mathbf{P}\{\mathbf{a_t} = \mathbf{e}_\ell\} &= \lambda_\ell \end{aligned} \tag{13}$$

where $\lambda_0 \in [0,1]$, and clearly $\sum_{\ell=1}^M \lambda_\ell = 1 - \lambda_0$ This means the observation vector likelihood $C_i(t) = \mathbf{P}\{Y_t | X_t = i\}$ can be expressed as the sum over disjoint events, with each event corresponding to an assignment of a particular observation y_t^ℓ to the target in some state, $X_t = i$, and all other observations false alarms from the sensor detector,

$$\begin{aligned} C_i(t) &= \mathbf{P}\{Y_t | X_t = i\} \\ &= \sum_{\ell=1}^M \mathbf{P}\{Y_t, \mathbf{a_t} = \mathbf{e}_\ell | X_t = i\} + \mathbf{P}\{Y_t, \mathbf{a_t} = \mathbf{0} | X_t = i\} \\ &= \sum_{\ell=1}^M \mathbf{P}\{\mathbf{a_t} = \mathbf{e}_\ell\} \mathbf{P}\{Y_t | X_t = i, \mathbf{a_t} = \mathbf{e}_\ell\} \\ &\qquad + \mathbf{P}\{\mathbf{a_t} = \mathbf{0}\} \mathbf{P}\{Y_t | \mathbf{a_t} = \mathbf{0}\} \\ &= \sum_{\ell=1}^M \lambda_\ell \mathbf{P}\{y_t^\ell | X_t = i, \mathbf{a_t} = \mathbf{e}_\ell\} \prod_{m \neq \ell}^M \mathbf{P}\{y_t^m | \mathbf{a_t} = \mathbf{e}_\ell\} \\ &\qquad + \lambda_0 \mathbf{P}\{Y_t | \mathbf{a_t} = \mathbf{0}\} \end{aligned} \tag{14}$$

In the first term of the final expression, the outer sum checks all possible assignments of detections to a target in state i, which are disjoint events. The inner product is simply a factorisation of $\mathbf{P}\{Y_t | X_t, \mathbf{a_t}\}$ due to the conditionally independent observations, given the association $\mathbf{a_t}$.

Analogously to the single observation model, for each y_t^m, we define a likelihood function given $\mathbf{a_t}$ as, $c_i^m(t) = \mathbf{P}\{y_t^m | X_t = i, \mathbf{a_t}\}$, $\forall\, i, m, t$. If $\mathbf{a_t} = \mathbf{e_m}$ then this likelihood is of the same form as in the previous section, since the association vector is pointing to observation y_t^m. The two other cases, when $\mathbf{a_t} \neq \mathbf{e_m}$ and $\mathbf{a_t} = \mathbf{0}$, we define as follows, assuming once again that non-target detections are generated in a similar way to target detections and each can be modelled via up to M independent clutter processes U_t

$$\begin{aligned} \mathbf{P}\{y_t^m | \mathbf{a_t} = \mathbf{e}_\ell\} &= \sum_{j=1}^N \mathbf{P}\{y_t^m | U_t = j, \mathbf{a_t} = \mathbf{e}_\ell\} \mathbf{P}\{U_t = j\} \\ &= \sum_{j=1}^N c_j^m(t) \mathbf{P}\{U_t = j\} \\ \mathbf{P}\{Y_t | \mathbf{a_t} = \mathbf{0}\} &= \prod_{m=1}^M \mathbf{P}\{y_t^m | \mathbf{a_t} = \mathbf{0}\} \\ &= \prod_{m=1}^M \sum_{j=1}^N \mathbf{P}\{y_t^m | U_t = j, \mathbf{a_t} = \mathbf{0}\} \mathbf{P}\{U_t = j\} \end{aligned}$$

Numerical examples of the multiple observation model above will be used in the numerical simulation section, with example realisations plotted. However, it is perhaps instructive to consider the noiseless case. Consider the case of 2 observations $Y_t = \{y_t^1, y_t^2\}$, this observation vector could correspond to

$$\begin{bmatrix} y_t^1 \\ y_t^2 \end{bmatrix} \rightarrow \begin{bmatrix} X_t = y_t^1 \\ U_t = y_t^2 \end{bmatrix} \textit{ or } \begin{bmatrix} U_t = y_t^1 \\ X_t = y_t^2 \end{bmatrix},$$

In this situation the underlying RC is still hidden, but only by ambiguity, since it is not known which observation in the list Y_t corresponds to the target.

IV. Optimal Filtering for Hidden Reciprocal Chains

In this section, we derive optimal filters for the estimation of the state sequence of a HRC, and for evaluation of the likelihood of a HRC observation sequence. Based on the construction of the observation models, the filter takes the same form regardless of the observation scenario, and so using the HRC filter in either simply becomes a matter of calculating and using the different observation likelihood functions. Therefore we present the HRC filter for a general observation sequence model.

Consider a HRC $\mathcal{Y}$ with state $\mathcal{X}$, known MB transition probability matrices $B_{i,j}^k(t)$, $t = 0, 2, \dots, T-1$, and known observation probabilities $C_i(t), t = 0, \dots, T, i = 1, \dots, N$. Given a sequence of observations, $Y_0, \dots, Y_T$, the aim here is to compute the filtered *a posteriori* probability (APP) mass functions

$$q_i(t) = \mathbf{P}\{X_t = i | Y_1, \dots, Y_t\} \ ,$$

for each $t = 0, \dots, T-1$, and each $i = 1, \dots, N$. These APPs can be calculated from the following joint probabilities, for each $k = 1, \dots, N$,

$$q_i(t) = \sum_{k=1}^N \mathbf{P}\{X_t = i, X_T = k | Y_1, \dots, Y_t\} = \sum_{k=1}^N q_i^k(t) \ ,$$

via the law of total probability. It has been shown that the joint process (X_t, X_T) is Markov [17], therefore analogously to the hidden Markov model filter (see e.g. [22]), it is easily

shown via Bayes' rule that $q_i^k(t)$ can be evaluated recursively, for $t = 1, \ldots, T-1$ by

$$q_i^k(t) = \frac{C_i(t) \sum_{j=1}^N B_{j,i}^k(t-1)\, q_j^k(t-1)}{h(t)} \ , \qquad (15)$$

where the normalising term is

$$\begin{aligned} h(t) &= \mathbf{P}\left\{Y_t | Y_0, \ldots, Y_{t-1}\right\} \\ &= \sum_{i,k=1}^N C_i(t) \sum_{j=1}^N B_{j,i}^k(t-1)\, q_j^k(t-1) \ . \end{aligned}$$

The overall computational cost of the above filtering recursions is $O(N^3 T)$. Initialisation at $t = 0$ is via

$$q_i^k(0) = \frac{C_0(i)\, \Pi_{i,k}}{h(0)},$$

$$\text{where} \quad h(0) = \mathbf{P}\left\{Y_0\right\} = \sum_{i,k=1}^N C_0(i)\, \Pi_{i,k}.$$

The observation log-likelihood can be evaluated via $\log \mathbf{P}\left\{Y_0, \ldots, Y_T\right\} = \sum_{t=0}^T \log h(t)$. This approach avoids any potential numerical underflow issues which might arise if the *conditional* (un-normalised) MB filters as defined in [16] were used. A track extraction type algorithm can be obtained by comparing the log sequence likelihood to a threshold.

A. Optimal Single and Multi-Observation HRC filters

As stated, in the case where the observations are generated via the single or multi-observation models of III-A, we simply use the correct corresponding observation likelihood, ie. (12) or (14), and substitute these likelihoods into the Markov bridge APPs of Equation (15).

Note that the complexity of this filter is no greater in the single observation scenario, however in the multi-observation case the complexity of this filter is $O(N^3 T + M^2 N^2 T)$, where the latter term is the burden of computing the likelihoods for the list of M observations at each time t. Thus if $M << N$, the complexity of the filter is dominated by the original term on the left, $O(N^3 T)$.

V. Track Extraction Detectors

The observation models we have presented, which represent the sensor level detection process, along with the associated filters, allow us to investigate the track extraction benefits gained by incorporating source-destination awareness into a detector's target model. We do this by defining several track extraction detectors, where each is a hypothesis test with two competing hypotheses. Based on the filters derived in the previous section we are able to formulate a likelihood ratio test in each scenario. Note that since an analytic form for the probability distribution of the sequence log likelihood under the null or alternative hypotheses can not be found, numerical approaches (such as simulations) are required in order to sensibly set the detection threshold.

The first class of detector we present attempts to decide whether a target exists in the cluttered environment. Thus in the single observation case the hypotheses are concerned with the value of ϵ in the mixture model. The null hypothesis is that the observations are all clutter generated ($\epsilon = 1$). The alternative hypothesis presumes there is a reciprocal target, and over many realisations we expect to receive target generated observations at a rate $(1 - \epsilon)$.

$$\begin{aligned} H_0 : &\quad \epsilon = 1 \quad \text{(no target present)} \\ H_R : &\quad \epsilon < 1 \quad \text{(a reciprocal target is present)} \end{aligned}$$

We call this a *reciprocal detector*. To compare the reciprocal model with a compatible Markov model, we form another detector where the alternative hypothesis instead has sequence log likelihood obtained from a standard HMC filter which assumes a target model with Markov dynamics.

$$\begin{aligned} H_0 : &\quad \epsilon = 1 \quad \text{(no target present)} \\ H_M : &\quad \epsilon < 1 \quad \text{(a Markov target is present)} \end{aligned}$$

The dynamics could be those of the base Markov process, or of the Schrödinger bridge with dynamics given by Equation (9). We call this detector a *Markov detector* if the base process is used to model dynamics, or a *Schrödinger detector* if a SB is used.

We also define corresponding detectors for the multi-observation case, which we parameterise via the priors over the association variable $\mathbf{a_t}$. Specifically, from (13) the null hypothesis will be parameterised as $\lambda_0 = 1$, and the alternative hypothesis as $\lambda_0 < 1$. The multi-observation likelihood for the null hypothesis is given by

$$\begin{aligned} \mathbf{P}\left\{\mathcal{Y} | H_0\right\} &= \prod_{t=0}^T \mathbf{P}\left\{Y_t | H_0\right\} \\ &= \prod_{t=0}^T \prod_{\ell=1}^M \sum_{i=1}^N \mathbf{P}\left\{y_t^\ell | U_t = i, \mathbf{a_t} = \mathbf{0}\right\} \mathbf{P}\left\{U_t = i\right\}. \end{aligned}$$

As before, we will define the temporally independent clutter process U_t to be uniform over the state space. In the case of noiseless observations, the observation sequence likelihood expression above simplifies further.

VI. Numerical Examples

In this section, we present numerical simulations to illustrate the performance of the proposed filtering algorithms for a HRC track extraction scheme. The simulated scenarios are designed to highlight the benefit that can be achieved in both state estimation and detection when source-destination awareness is incorporated into the target model. We compare the HRC tracker (estimator and detector) to existing Hidden Markov chain (HMC) trackers similar to the track extraction models of [1], [2], and also compare to a tracker based on a Schrödinger bridge, which we call a Hidden Schrödinger Chain (HSC) tracker.

The HMC tracker uses the base process defined in subsection II-B as its target model, and the HRC and HSC trackers use target models with dynamics derived from this same base process. The models differ in terms of the statistical characterisation of the end-points and their dynamics ; the

HRC has a specified joint distribution Π on the end-points, the HMC has only initial state distribution π_0, being the marginal derived from Π ; and the HSC has initial and final distributions equal to the marginal distributions of Π.

In all simulations the target trajectory data is generated according to the RC target model, for different distributions Π, and thus the HMC and HSC filters and detectors are mismatched while the HRC is matched. This is in line with the premise of the paper, that targets on courser time scales have at least a simple notion of intent, in this case source-destination awareness.

Firstly, in subsection VI-A we briefly describe the simulation environment. In VI-B we describe and justify the kinds of source-destination awareness we have chosen to encode with the RC model. In VI-C, we present the detection and filtering results, for both single and multiple observation scenarios, before we discuss these results in VI-D.

A. Simulation Environment

In all the examples presented, the states take values on a regular two-dimensional 8×8 cellular gridworld - the *region of interest* (ROI). We restrict our RCs to those whose local dynamics are built from one-step Markovian random walks, parameterised by the probability of remaining in a cell p_R and equal probabilities of moving to the neighbouring cells (states). Note that neighbouring cells include those on the diagonal, meaning the random walk is '8-connected', rather than 4-connected, though this is an arbitrary choice. Jumps outside the ROI are not permitted. Fig. 1 illustrates the basic nature of the state space and some sample trajectories.

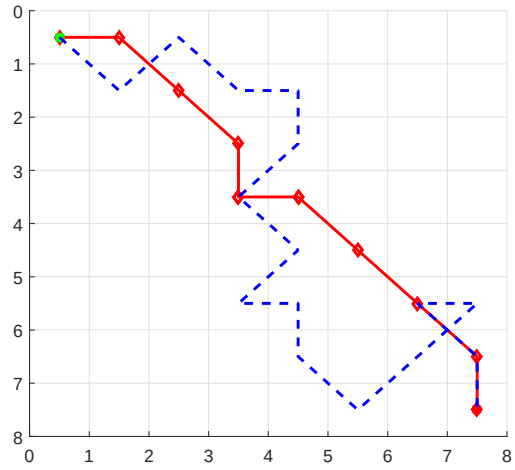

Fig. 1. *Simple source destination aware trajectories*: The paths of 2 RC targets which cross the 8×8 lattice from $(1,1)$ to $(8,8)$. Green markers denote the origin state of any realisation, and the red denote the end state. Each trajectory corresponds to that of a Markov bridge with a fixed start, the simplest type of destination awareness. The red target has $T = 12$ steps to reach state $(8,8)$, while the blue has $T = 32$.

The trajectories are realisations of two different RCs. Note that the plotted paths do not show if a target remains in or revisits a particular state. In order to generate an observation sequence, the independent clutter process U_t is also realised. A sensor detection in the observation sequence is obtained by adding zero-mean Gaussian noise to the centre coordinates of the cell that the target or clutter is in at time t. Note this choice is illustrative and easily be generalised. In the single observation model, whether the target or clutter selected of course depends of course on the clutter rate ϵ, and similarly in the multi-observation case. The added Gaussian noise has equal variance σ^2 in the x and y directions with the x and y components of the noise being statistically independent. Fig. 4 presents an RC and the observation process in the single observation scenario.

B. Joint Endpoint Distributions

The performance of all of the detectors and filters considered depends on the ability of the underlying target model to accurately describe the target's dynamics. As stated, if there is a target, it is reciprocal. It would be expected that the HRC based detector should perform the best since it is matched to the data, but the performance benefit may be small, depending on how the destination awareness affects the RC dynamics, which are built from, and therefore different to, those of the base process. Therefore our numerical studies have focused on the parameters that determine the difference between the RC dynamics and the base process dynamics.

The red path, in Fig. 1, shows a target moving a distance of 7 steps in an interval length of $T = 12$, whereas the blue target has $T = 32$ to move the same distance. Both targets are realisations from RCs with the same Π and dynamics built from the same MC dynamics A, but their trajectories are very different. The red target is more restricted and moves with a clearer purpose, while the blue target appears to move more randomly, with dynamics closer to the original Markov dynamics.

In general, it is not just the sequence length T that determines how restricted or random the average dynamics of a RC target will be. This will also depend on a) the size of the state space and the ease with which a target can traverse the lattice (whether it can only move one step or more, and with what likelihoods), which are both encoded into the transition matrix A, and in terms of average dynamics b) the distribution over the source-destination pairs, Π. Since for these simulations the state space and the transition matrix A are fixed (and restricted to one step dynamics), we propose an empirical formula that relates the extent to which the RC dynamics are less random than the MC dynamics, to the free parameters: the interval length T, the minimum number of steps $d_{min}(i,j)$ between source-destination pairs (i,j), and the distribution over pairs Π.

$$\beta = \frac{\sum_{i,j=1}^{N} d_{min}(i,j)\Pi(i,j)}{T} \in [0,1] \tag{16}$$

Based on intuition and general observation from simulations, we claim that the term β will in turn predict the benefit that HRC will provide over Markov trackers according to a linear relationship, with a constant offset κ,

$$\text{HRC Benefit} \propto \beta + \kappa \tag{17}$$

We will refer to β as the benefit indicator, and define benefit in the following subsection, in the detection case. Notionally if β is closer to 0 we expect the RC dynamics to be similar to the MC dynamics, whereas if β is unity, the RC dynamics result in a deterministic path and thus would appear quite different to the MC dynamics. We will present results testing this claim in the following subsection.

As an example, consider a RC model which only allows targets to start in the corners and cross the 8×8 state space to its opposite corner, as in figure 2. We will call endpoint distributions of this form a crossing RC (CRC). If we then set $T = 8$, the targets then have the minimum time required to move the $d_{min}(i, j) = 7$ steps between its start and final state. This motion is deterministic according to the RC, and while possible according to the MC, highly unlikely. The deviation for this case is $\beta = 1$. We would expect HRC trackers to perform best for target behaviour of this type. Another interesting test case is for when the target is still non-Markov but the benefit indicator β is very low, or 0. An example of this case is the loitering RC , for which the indicator is indeed $\beta = 0$, since the minimum number of steps $d_{min}(i, j) = 0$ (as $i = j$ for all source-destination pairs, see Eq. (8)).

We have chosen joint endpoint distributions Π which range between the two cases given, crossing and loitering, according to the following construction

$$\text{System } \Pi = \alpha\Pi_{CRC} + (1 - \alpha)\Pi_{LRC} \ , \tag{18}$$

where $\alpha \in [0, 1]$, and LRC denotes a loitering RC as before. We have used this construction to study the performance of the filters and detectors more systematically. Thus, as a consequence of the claim in equation (16), it should be expected that the performance benefit can be controlled by the sequence length T and the parameter α, which ranges from what we expect to be best case ($\alpha = 1$) to worst case ($\alpha = 0$). We can express the benefit indicator as a function of α,

$$\beta = \frac{\alpha}{T}\, d_{CRC} \tag{19}$$

since $d_{min} = 0$ for all loitering source-destination pairs, and assuming the same, constant number of steps d_{CRC} for all crossing source-destination pairs. In the 8×8 state space $d_{CRC} = 7$ steps. Note that if $\alpha = 0$, and all trajectories are loitering, T can be lowered such that the end state is always reachable for any T.

C. Results

In this subsection we present results for both the detection and filtering algorithms in the simulation environment described. All three trackers are applied in the single and multi-observation scenarios, with the results consistent across both scenarios. We begin first with results on detection, specifically with a focus on testing the claims of Equations (16) and (17), when parameterised by α as in Equation (18). We then present results on the estimation performance of the filters. Despite track extraction being largely about detection, the filtering results give us an insight into the detection mechanism, and why the different models perform differently.

Results on the effect of clutter and general measurement noise reflect the expected property that the benefit of better models increases with worsening tracking conditions. We provide evidence of this by comparing HMC to HRC performance as conditions worsen, strengthening the general argument made in the paper. We leave A constant as mentioned for the primary detection and filtering simulations with the probability of the remaining in a cell fixed to $p_R = 0.5$. Following the primary results we will vary p_R to investigate its effect on performance. Note that unless otherwise stated, results are obtained from $10,000$ independent realisations.

Detection:

Detection results are presented using receiver operator characteristic (ROC) curves to begin with in Figure 2, which plot the estimated probability of detection against the estimated false alarm rate as the detection threshold τ is varied. We simulate realisations of both hypotheses in equal number, RC target present and no target present, reflecting the uniform priors chosen. The detectors perform the likelihood ratio test using the uniform priors, equal penalties for incorrect decisions and no penalties for correct decisions, that is, a minimum error probability test.

To begin with we consider the single observation scenario with a clutter rate $\epsilon = 50\%$ and increase α from 0 to 1. This means we should expect to see the benefit of the RC increase with α, which is shown in Fig. 2. The multi-observation case obtains similar results across all α settings.

In Figure 3 we present results testing the validity of the proposed relationship and formula in (17), with the area under the curve (AUC) of the ROC curve taken to be the measure of benefit that HRC provides over the HMC tracker. The AUC is a valid measure of aggregate detector performance provided that one of the detectors being compared has ROC curve consistently above the other curve [23], which is the case here. The benefit indicator β is varied via (19) as α increases from 0 to 1, for different fixed interval lengths T, corresponding to each curve. We can see that the proportional relationship is approximately correct, especially as sequence length T increases. We see also that the curves become shorter, which is to be expected according to Equations (16) and (17).

Filtering:

In the filtering example of Figure 4, we set $\alpha = 1$ and purely considers CRC targets, with a path length of $T = 16$, in the single observation scenario. The clutter rate was set to 25% ($\epsilon = 0.25$), meaning that on average 75% of any observation sequence will be of target generated measurements.

The performance of the filters is measured by the RMSE of the online conditional means. The conditional mean is a continuous real valued estimate corresponding to the $2D$ gridworld state space,

$$\hat{X}_t = \mathbf{E}\left\{X_t | Y_0, \ldots Y_t\right\}$$

Given $X_t^r, Y_t^r, \hat{X}_t^r$ are the states, observations and conditional mean estimates corresponding to a particular realisation $r \in \{1, \ldots, R\}$ at each time t, the RMSE of the conditional mean estimates is calculated using the Euclidean distance from the

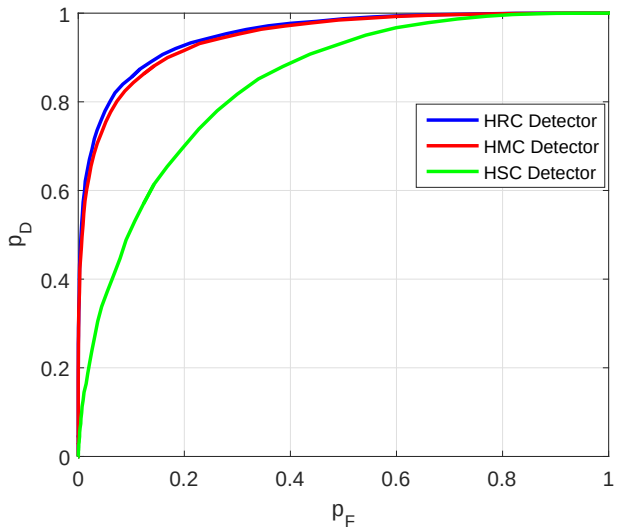

(a)

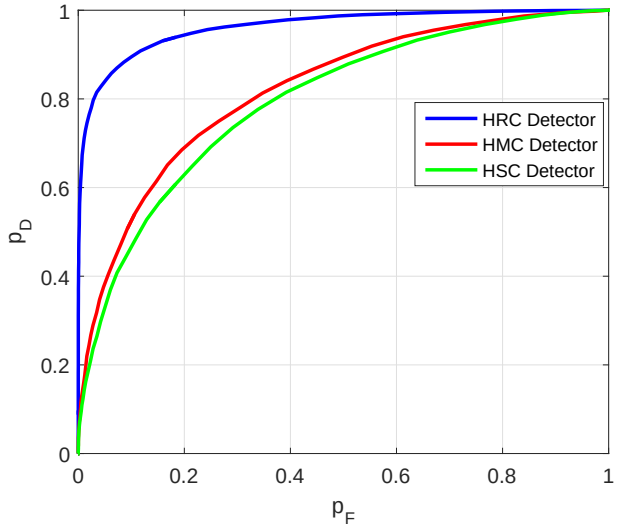

(b)

Fig. 2. *Single Observation Detection Results*: ROC curve for a single observation system with clutter rate $\epsilon = 50\%$ and sensor detection noise $\sigma^2 = 1$. The RC has interval length $T = 16$, and in a) $\alpha = 0$, b) $\alpha = 1$.

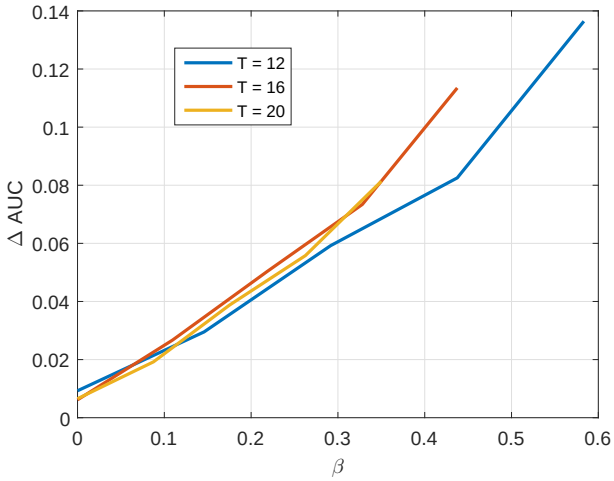

Fig. 3. ΔAUC *vs.* β: Benefit provided by RC taken to be area between ROC curves oh HRC and HMC (ΔAUC) in a single observation scenario, with $\epsilon = 0.4$, $\sigma^2 = 1$, $T = 16$ as α increases.

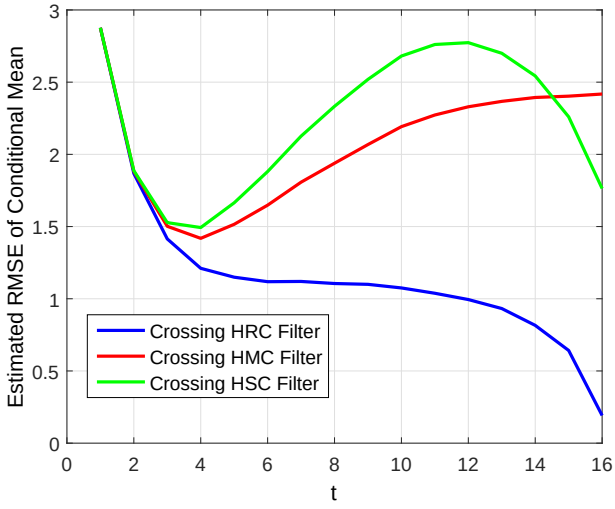

Fig. 4. *Single Observation Filtering Results*: Estimated RMSE of the Conditional Mean for each of the trackers with the clutter rate of $\epsilon = 0.25$, and sensor detector noise of $\sigma^2 = 1$. The RC has interval length $T = 16$,, and $\alpha = 1$.

target's state coordinates, x_t^r, y_t^r to the estimate coordinates, $\hat{x}_t^r, \hat{y}_t^r$,

$$\text{RMSE}_{CM}(t) = \sqrt{\sum_{r=1}^{R} \frac{(\hat{x}_t^r - x_t^r)^2 + (\hat{y}_t^r - y_t^r)^2}{R}}$$

where. The RMSE of the conditional mean provides a good measure of tracking accuracy. The results for the multi-observation scenario are very similar for online estimation, to within a small scaling factor.

As stated, we expect that HRC performance benefit increases with worsening tracking conditions, since the model information becomes more useful. Figure 5 shows the difference in average per sample RMSE between the HMC and HRC filters, in the multi-observation scenario between, as the number of observations per time increases. We define the average per sample (APS) RMSE of the filters as,

$$\text{RMSE}_{APS} = \sum_{r=1}^{R} \frac{1}{R} \sqrt{\sum_{t=1}^{T} \frac{(\hat{x}_t^r - x_t^r)^2 + (\hat{y}_t^r - y_t^r)^2}{T}}$$

This produces a similar measure to the RMSE defined previously, but across the entire sequence, so that we may describe performance as the interval length T is varied. Each curve in Figure 5 represents a different set of independent realisations corresponding to the number of observations at each time, increasing from $M = 2$ to 6, when all targets are crossing. We set the *a priori* probability of no target detection, to $\lambda_0 = 0$ (see Equation (13)), since the effect of $\lambda_0 > 0$ is seen

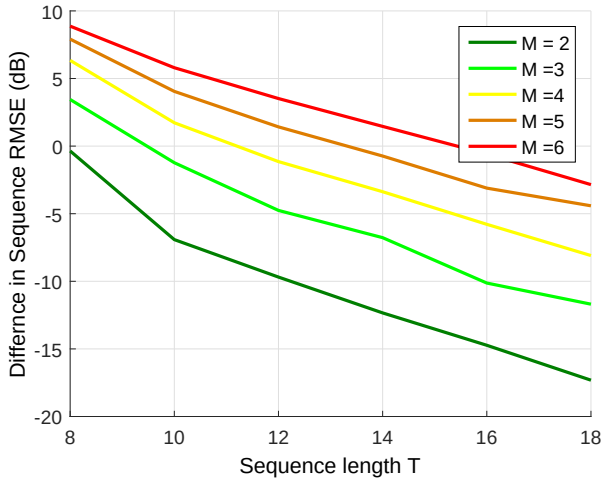

Fig. 5. *Multi Observation Filtering Results*: Filtering performance difference of HRC over HMC model. The vertical axis is the benefit in RMSE in dB of the average per sample RMSE as interval length T is varied, for different number of observations per time M, from $M = 2$ (dark green) to $M = 6$ (red). The horizontal axis is the sequence length T.

in the single observation scenario, thus the multi-observation filter is resolving ambiguity. The phenomena of increasing performance of HRC over alternative Markov models was observed across all worsening conditions; increased sensor detector noise σ^2, clutter rate ϵ as well as the number of targets M. We note also that the benefit of HRC over HMC decreases with increasing sequence length T.

Varying transition probabilities of A:

The final numerical simulations investigate the effect that varying the one-step transition probabilities A of the base process have on the HRC and HMC trackers' performance, using RMSE of the conditional mean estimates to measure the performance once more. We omit the HSC tracker for reasons that will become clear in the discussion. Figure 6 shows the performance of the HRC and HMC filters respectively, in the single observation scenario, with only CRC targets ($\alpha = 1$) as the probability of remaining in a vertex p_R is varied from 0.2 to 0.8. We see that the HRC is largely invariant to the variation in p_R, while the HMC is not, estimating more poorly for higher p_R. Figure 7 however, shows no difference between the HRC and HMC filtering performance when all the targets loiter.

This effect of p_R on the HRC and HMC trackers does not carry through from filtering to detection. We see this in Figure 8, where the difference in AUC between the HMC and HRC ROC curves is not effected by the change in p_R.

D. Discussion

The results show that the HRC tracker, which is matched to the model generating realisations, performs best when compared to the Markov trackers, as expected. Furthermore the proposed formula (16) is supported, since the claim of linear proportionality (17) appears to be approximately correct, especially for longer T. Despite lacking a direct probabilistic

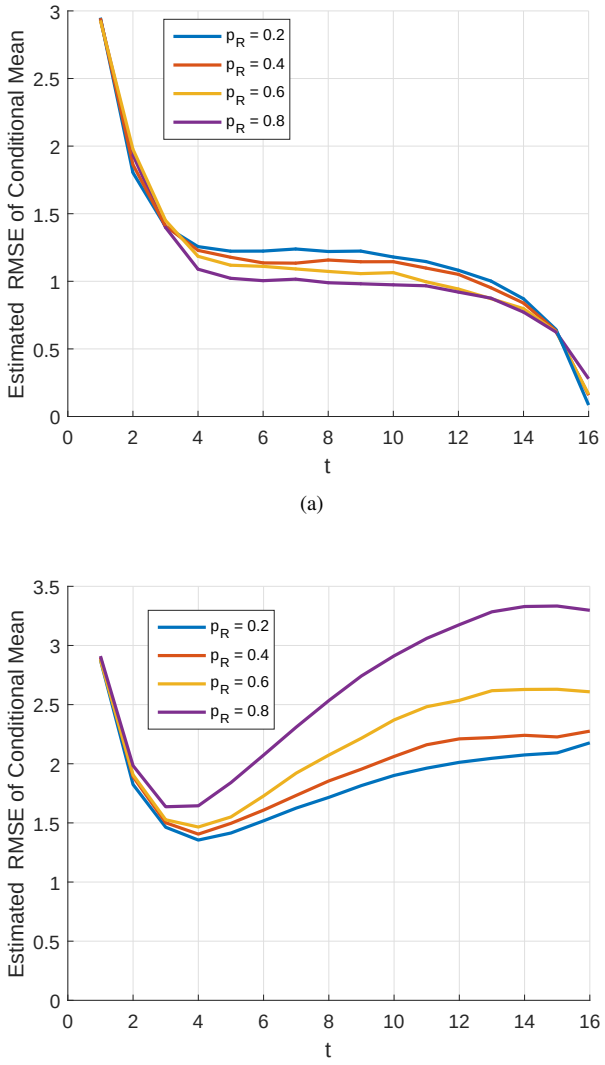

Fig. 6. *Filtering conditional mean RMSE for different p_R*: target is a crossing RC ($\alpha = 1$), $\epsilon = 0.5$, $T = 16$, a) HRC b) HMC.

interpretation, Equation (16) provides a good intuitive description of the relationship between the parameters of the HRC model and the benefit over the HMC model.

When we consider the HMC and HSC results specifically we discover that although the Schrödinger bridge model incorporates more future information than the Markov model, it performs worse. The first reason for this is that the Schrödinger bridge construction introduces source destination pairs not present in the original Π, as indicated in section II-B. This can be understood by considering the endpoint marginals obtained from the RC and passed to the SB. In the loitering case for example if the RC target can loiter anywhere uniformly in the state space, this will result in uniform marginals across both start and end points for the SB. Thus the HSC tracker will generally 'expect' trajectories that will never be realised from the RC model. The second reason for these errors is that the clutter observation model allows the HSC tracker to assign the

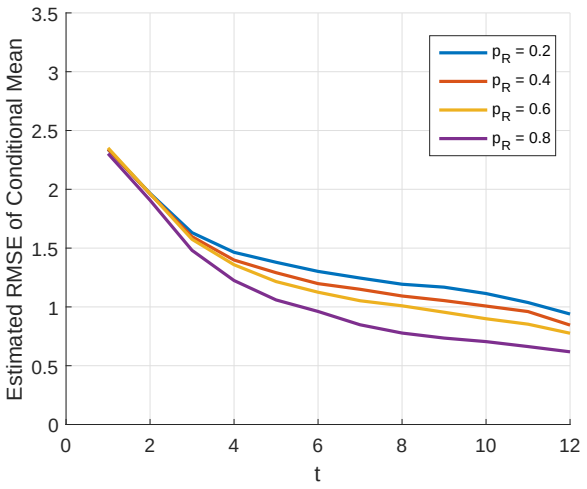

(a)

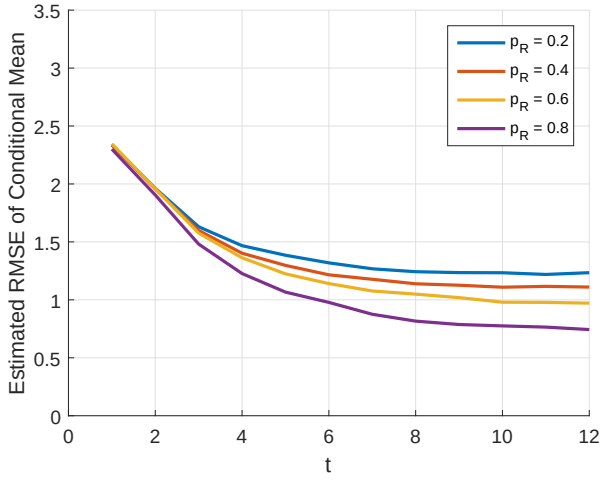

(b)

Fig. 7. *Filtering conditional mean RMSE for different* p_R: target is a loitering RC ($\alpha = 0$), $\epsilon = 0.5$, $T = 12$, a) HRC b) HMC.

majority of measurements to clutter. This effect is evident in both the detection and filtering results, such as Figures 2(b) and 4, and points to the advantage of RC over SB modelling of targets with future information.

We can explain the HSC results further with a specific example of a crossing target in the single observation scenario. Figure 9 is a plot of one particularly poor performance of the HSC tracker. The figure displays the target path $\mathcal{X}$, a noiseless path of the sensor detections (both target and clutter), the observations $\mathcal{Y}$ and the HSC filter's maximum *a posteriori* probability (MAP) estimates. Note there are 12 HSC filter estimates (green), however many overlap. The majority of the HSC estimates are focused around the starting vertex $(1, 1)$, since under its model it is possible for a target to start and end at the same location (due to the endpoint marginals). However, at time $t = 9$ the filter 'realises' the measurement sequence makes more sense with a trajectory of $(1, 1)$ to $(8, 8)$. This effect corresponds to Figure 3, where the RMSE in the HSC

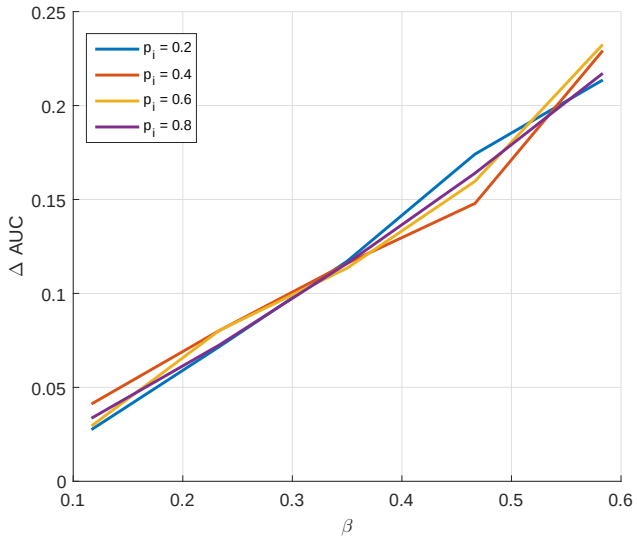

Fig. 8. ΔAUC *vs.* β *for different* p_R: Benefit once more taken to be the difference in area under the ROC curves of the HRC and HMC trackers (ΔAUC). $T = 16$ as α is varied, $\epsilon = 0.5$. Graph shows invariance with respect to p_R.

conditional mean estimates decreases close to T.

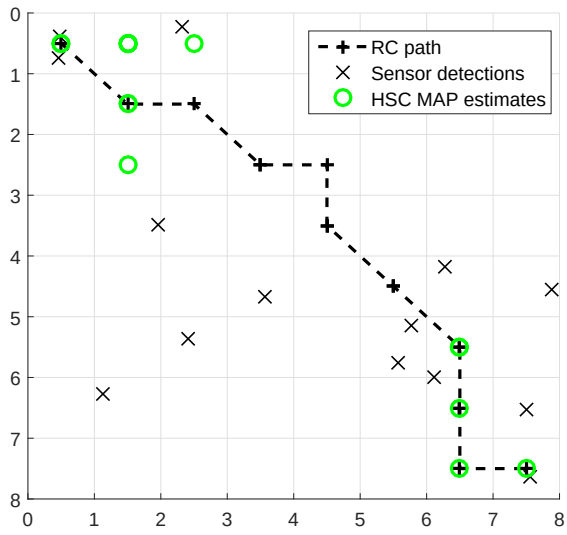

Fig. 9. *HSC estimates predicting incorrect trajectory*: The figure above shows the target's path from $(1, 1)$ to $(8, 8)$ (dotted black line), the sensor detections (black points) and the HSC filter MAP estimates (green circles), for a crossing RC ($\alpha = 1$) with $T = 16$. Note there are in fact 16 points for each of the plots, however some overlap. It can be seen the HSC predicts a path close to the origin (state 1, $(1, 1)$) until $t = 9$ when the last 4 data points are much closer to the target path.

Finally, we consider Figures 6 and 7, which present RMSE of the conditional mean of the filters as p_R varies. Since a MC will have a distribution over its state at time t (i.e. $\pi_0(A)^t$) that is more concentrated about its origin for higher p_R, it makes sense that for higher p_R the HMC should produce

worse estimates for a CRC but better estimates for a LRC. By their definitions a CRC forces the target to move the furthest distance in time T whereas a loitering RC alters the base process dynamics only to return to the origin with probability 1, which explains the HMC performance. As state before however, this effect does not carry over to the detector performance.

Thus the two main results from the numerical studies are, primarily that the benefit of the HRC is controlled by the parameters of the RC model, and secondly that the HSC is a poor model for tracking targets that proceed from an initial state to arbitrary final state, and an analyst would be better served with a simple HMC tracker that uses the globally naive base process.

VII. Conclusion

This paper has considered the problem of how target models which incorporate a simple notion of intent can improve the tracking performance over alternative Markov target models. To do this, track extraction scenarios were recreated by defining two observation models of detections generated by thresholding sensor returns, with targets being modelled as reciprocal chains and clutter as a temporally uncorrelated process. A reciprocal chain can be constructed from a MC base process together with a joint endpoints distribution which encodes the source-destination awareness of a target. Normalised HRC filters were presented for the observation models introduced and the normalising constants were used to form likelihood based detectors. These HRC filters and detectors were compared via numerical simulations to both a HMC tracker with the base process as its target model, and a Markov tracker based on the Schrödinger bridge, which can encode some future information.

Numerical simulation design was of a HRC on a 2D cellular gridworld of fixed size, with states as cells, and dynamics restricted to one-step walks over cells. The central result from these numerical simulations was that the benefit of the source-destination awareness of an RC, built from a base process with known one-step Markov dynamics, can be related to the sequence length, the distance between source and destination and the likelihood of that a target with that source and destination being realised. Specifically, if benefit is taken to be the difference in area under the ROC curves of the reciprocal and base Markov detector, it was found to vary approximately linearly with a simple function of the RC parameters in a specific scenario.

An important secondary result is that the Schrödinger bridge tracker, despite having the correct start and endpoint distributions pre-specified, in fact performs worse than even a Markov tracker using base process.

Research tasks that follow directly from this work are first and foremost a generalisation of the central result of the simulations, with fewer restrictions on the base process (ie. allowing more than simple one-step dynamics) and subsequently replacing the empirical benefit indicator β with a more general metric to indicate how the reciprocal dynamics deviate from Markov dynamics as a function of the RC parameters. Most pressing research would be to address the problem of parameter estimation, and the related questions of how suitable the algorithms presented are in a real tracking system, and whether HRCs are strictly batch methods or could be feasibly employed online. Of the many parameters to be estimated, the most interesting and challenging is the joint endpoints distribution over the source and destinations, and how to select an appropriate corresponding T. In a real tracking environment there would be many targets behaving in various ways, with the potential for different RC targets existing at once. This research would benefit greatly from considering situations with corresponding real data, such as in [5]-[8] and the references therein.

Acknowledgements

The support of the Defence Science and Technology Group, Australia, and the Data to Decisions CRC is gratefully acknowledged. We would also like to thank Dr. Jason Williams for his useful insight and contribution.

References

[1] G. van Keuk , "Sequential track extraction", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34 , no. 4, , pp. 1135 - 1148, Oct. 1998.

[2] M. Wieneke , W. Koch, "On sequential track extraction within the PMHT framework", *EURASIP Journal on Advances in Signal Processing*, pp. 1-13, Jan. 2008.

[3] S. Blackman, "Multiple Hypothesis Tracking for Multiple Target Tracking", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 19 , no. 1, Jan. 1998.

[4] S. Blackman, R. Popoli *Multiple Design and Analysis of Modern Tracking Systems*, Norwood, MA: Artech House, 1999.

[5] D. Patterson, L. Liao, d. Fox, H. Kautz, "Inferring High-Level Behavior from Low-Level Sensors", *Proc. 5th Int. Conf. Ubiquitous Computing*, LNCS 2864, Springer-Verlag, pp. 73-89, 2003.

[6] M. Morelande, M. Duckham, A. Kealy, J. Legg "Bayesian path estimation using the spatial attributes of a road network", *Proc. IEEE ICASSP*, pp. 4090-4094, April 2015.

[7] J. Krumm, E. Horvitz, "Predestination: Inferring destinations from partial trajectories", *Proc. 8th Int. Conf. Ubiquitous Computing*, pp. 243-260, 2006.

[8] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning", *Proc. AAAI*, pp. 1433-1439, 2008.

[9] D.A. Castanon, B.C. Levy, and A.S. Willsky, "Algorithms for the incorporation of predictive information in surveillance theory", *Int. J. Systems Sci.*, vol. 16, 1985, pp. 367-382.

[10] J.E. Kulkarni and L. Paninski, "State-space decoding of goal-directed movements," *IEEE Signal Processing Magazine*, vol. 25, pp. 78, 2008.

[11] B. Jamison, "The Markov Processes of Schrödinger," *Probability Theory and Related Fields*, vol. 32, no. 4, pp. 323–331, 1975.

[12] M. Pavon and F. Ticozzi, "Discrete-Time Classical and Quantum Markovian Evolutions: Maximum Entropy Problems on Path Space," *Journal of Mathematical Physics*, vol. 51, no. 4, 2010.

[13] T. T. Georgiou and M. Pavon, "Positive Contraction Mappings for Classical and Quantum Schrödinger systems," *Journal of Mathematical Physics*, vol. 56, no. 3, p. 033301, 2015.

[14] B. Jamison, "Reciprocal Processes", *Z. Warscheinlichkeitstheorie verw. Gebiete*, 30, pp. 65-86, 1974.

[15] B. C. Levy, R. Frezza and A. J. Krener, "Modeling and Estimation of Discrete-Time Gaussian Reciprocal Processes", *IEEE Trans. Automatic Control*, v. 35, no. 9, Sept. 1990, pp. 1013-1023.

[16] L. B. White and F. Carravetta, "Optimal smoothing for finite state hidden reciprocal processes", *IEEE Trans. Automatic Control*, v. 56, no. 9, pp. 2156-2161, 2011.

[17] L. B. White and F. Carravetta, "New Normalised Bayesian Smoothers for Signals Modelled by Non-causal Compositions of Reciprocal Chains", *Proc. IEEE Statistical Signal Processing Workshop*, Gold Coast, pp. 205-208, 2014.
[18] L. B. White and H. Vu, "Maximum Likelihood Sequence Estimation for Hidden Reciprocal Processes", *IEEE Trans. Automatic Control*, vol. 58, no. 10, pp. 2670-2674, Oct. 2013.
[19] M. Fanaswala and V. Krishnamurthy, "Detection of Anomalous Trajectory Patterns in Target Tracking via Stochastic Context-Free Grammars and Reciprocal Process Models", *IEEE J. Sel. Topics in Signal Processing*, v.7, no. 1, pp. 76-90, 2013.
[20] M. Fanaswala, V. Krishnamurthy and L. B. White, "Destination-Aware Target Tracking via Syntactic Signal processing", *Proc. IEEE Int. Conf. Acoustics, Speech and Signal processing*, Prague, pp. 3692-3695, May 2011.
[21] G. Stamatescu, A. Dick and L. B. White "Multi-camera tracking of intelligent targets with Hidden Reciprocal Chains", *Digital Image Computing: Techniques and Applications*, 2015. (DOI: 10.1109/DICTA.2015.7371287)
[22] Y. Ephraim and N. Merhav, "Hidden Markov Processes", *IEEE Trans. Information Theory*, v. 48, no. 6, pp. 1518-1569, 2002.
[23] D. J. Hand, "Measuring classifier performance: a coherent alternative to the area under the ROC curve", *Mach Learn*, vol. 77 , no. 1 , pp.103 -123 2009.

George Stamatescu received the degrees of B. E. (hons) and B. Maths & Comp. Sc. in 2013 from the University of Adelaide, where he is currently a graduate student with the School of Electrical and Electronic Engineering. His research interests lie in the application of finite state stochastic processes in tracking and artificial intelligence problems.

Langford B white obtained the degrees of B. Sc. (mathematics), B. E. (hons) and Ph. D. degrees in Electrical Engineering in 1984, 85 and 89 from the University of Queensland. From 1986-1999 he worked for the Defence Science and Technology Organisation in South Australia. Since 1999 he has been a Professor in the School of Electrical and Electronic Engineering at the University of Adelaide. Prof Whites research interests include signal processing, control and game theory.

Riley Bruce-Doust received the degrees of B. E. (hons) and B. Sc. in 2009 from the Australian National University. He is currently a graduate student with the School of Electrical and Electronic Engineering at the University of Adelaide. He is researching properties of finite-state stochastic processes.